

Exact symbolic artificial intelligence for faster, better assessment of AI fairness

August 9 2021, by Rachel Paiste



Credit: Pixabay/CC0 Public Domain

The justice system, banks, and private companies use algorithms to make decisions that have profound impacts on people's lives. Unfortunately, those algorithms are sometimes biased—disproportionately impacting people of color as well as individuals in lower income classes when they apply for loans or jobs, or even when courts decide what bail should be set while a person awaits

trial.

MIT researchers have developed a new artificial intelligence programming [language](#) that can assess the fairness of algorithms more exactly, and more quickly, than available alternatives.

Their Sum-Product Probabilistic Language (SPPL) is a [probabilistic programming](#) system. Probabilistic programming is an emerging field at the intersection of programming languages and artificial intelligence that aims to make AI systems much easier to develop, with early successes in computer vision, common-sense data cleaning, and automated data modeling. Probabilistic programming languages make it much easier for programmers to define probabilistic models and carry out probabilistic [inference](#)—that is, work backward to infer probable explanations for observed data.

"There are previous systems that can solve various fairness questions. Our system is not the first; but because our system is specialized and optimized for a certain class of models, it can deliver solutions thousands of times faster," says Feras Saad, a Ph.D. student in electrical engineering and computer science (EECS) and first author on a recent paper describing the work. Saad adds that the speedups are not insignificant: The system can be up to 3,000 times faster than previous approaches.

SPPL gives fast, exact solutions to probabilistic inference questions such as "How likely is the model to recommend a loan to someone over age 40?" or "Generate 1,000 synthetic loan applicants, all under age 30, whose loans will be approved." These inference results are based on SPPL programs that encode probabilistic models of what kinds of applicants are likely, a priori, and also how to classify them. Fairness questions that SPPL can answer include "Is there a difference between the probability of recommending a loan to an immigrant and

nonimmigrant applicant with the same socioeconomic status?" or "What's the probability of a hire, given that the candidate is qualified for the job and from an underrepresented group?"

SPPL is different from most probabilistic programming languages, as SPPL only allows users to write probabilistic programs for which it can automatically deliver exact probabilistic inference results. SPPL also makes it possible for users to check how fast inference will be, and therefore avoid writing slow programs. In contrast, other probabilistic programming languages such as Gen and Pyro allow users to write down probabilistic programs where the only known ways to do inference are approximate—that is, the results include errors whose nature and magnitude can be hard to characterize.

Error from approximate probabilistic inference is tolerable in many AI applications. But it is undesirable to have inference errors corrupting results in socially impactful applications of AI, such as automated decision-making, and especially in fairness analysis.

Jean-Baptiste Tristan, associate professor at Boston College and former research scientist at Oracle Labs, who was not involved in the new research, says, "I've worked on fairness analysis in academia and in real-world, large-scale industry settings. SPPL offers improved flexibility and trustworthiness over other PPLs on this challenging and important class of problems due to the expressiveness of the language, its precise and simple semantics, and the speed and soundness of the exact symbolic inference engine."

SPPL avoids errors by restricting to a carefully designed class of models that still includes a broad class of AI algorithms, including the decision tree classifiers that are widely used for algorithmic decision-making. SPPL works by compiling probabilistic programs into a specialized data structure called a "sum-product expression." SPPL further builds on the

emerging theme of using probabilistic circuits as a representation that enables efficient probabilistic inference. This approach extends prior work on sum-product networks to models and queries expressed via a probabilistic programming language. However, Saad notes that this approach comes with limitations: "SPPL is substantially faster for analyzing the fairness of a decision tree, for example, but it can't analyze models like neural networks. Other systems can analyze both neural networks and decision trees, but they tend to be slower and give inexact answers."

"SPPL shows that exact probabilistic inference is practical, not just theoretically possible, for a broad class of probabilistic programs," says Vikash Mansinghka, an MIT principal research scientist and senior author on the paper. "In my lab, we've seen symbolic inference driving speed and accuracy improvements in other inference tasks that we previously approached via approximate Monte Carlo and deep learning algorithms. We've also been applying SPPL to probabilistic programs learned from real-world databases, to quantify the probability of rare events, generate synthetic proxy data given constraints, and automatically screen data for probable anomalies."

The new SPPL probabilistic programming language was presented in June at the ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI), in a paper that Saad co-authored with MIT EECS Professor Martin Rinard and Mansinghka. SPPL is implemented in Python and is available open source.

More information: Feras A. Saad et al, SPPL: probabilistic programming with fast exact symbolic inference, *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (2021). [DOI: 10.1145/3453483.3454078](https://doi.org/10.1145/3453483.3454078)

Provided by Massachusetts Institute of Technology

Citation: Exact symbolic artificial intelligence for faster, better assessment of AI fairness (2021, August 9) retrieved 23 April 2024 from <https://techxplore.com/news/2021-08-exact-artificial-intelligence-faster-ai.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.